

HIGH FREQUENCY BUS SYSTEM

Donald V. Perino, a citizen of the United States of America, residing at 1630 William Henry Court, Los Altos, California 94024;

Billy Wayne Garrett, Jr., a citizen of the United States of America, residing at 928 Wright Avenue, #405, Mountain View, California 94043;

Haw-Jyh Liaw, a citizen of Taiwan, residing at
120 E. Remington Drive, #102, Sunnyvale,
California 94087;

David Nguyen, a citizen of the United States of America, residing at 5315 Frachini Circle, San Jose, California 95136;

Srinivas Nimmagadda, a citizen of India,
residing at 1050 Benton Street, Apt. 3202,
Santa Clara, California 95050;

James A. Gasbarro, a citizen of the United States of America, residing at 1613 Notre Dame Drive, Mountain View, California 94040; and

Richard DeWitt Crisp, a citizen of the United States of America, residing at 1052 Rickenbacker Street, San Jose, California 95128.

Rambus, Inc.

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
(415) 576-0200

HIGH FREQUENCY BUS SYSTEM

5

BACKGROUND OF THE INVENTION

As computer processors increase in speed they require increased information bandwidth from other subsystems supporting the processor. An example is the large amount of bandwidth needed by video and 3D image processing from a computer memory subsystem. Another example is a main memory subsystem. One or more high frequency buses are typically employed to provide the bandwidth required. The higher the frequency of operation of the bus, the greater the requirement that the signals on the bus have high-fidelity and equal propagation times to the devices making up the subsystem. High-fidelity signals are signals having little or no ringing and controlled and steady rising and falling edge rates. Many obstacles are encountered in assuring the uniform arrival times of high-fidelity signals to devices on the bus. One such obstacle is a requirement that a subsystem be modular, meaning that portions of a subsystem may be added and possibly removed. The modularity requirement implies that devices that are part of the modular subsystem be mounted on a separate substrate or module which couples to another board, the motherboard. It also implies the use of connectors if both addition and removal is required. Other obstacles are the number of layers of the motherboard on which routing of the bus is allowed and whether the bus is routed in a straight line or routed with turns. Too few layers on a motherboard, or a module, and turns of the lines may not permit the construction of the bus lines in a way necessary to achieve uniform arrival times of high-fidelity signals to devices on the bus.

Modular subsystems in computers have numerous advantages, some of which are field upgradability, replacement of a failing device, flexibility of initial configuration, and increased device density. Currently, so called SIMMs (single in-line memory module) and DIMMs (dual in-line memory modules)

5

10

30

35

tapped, it is often necessary to insert a resistor 320, as shown in Fig. 3, in series with the stub to mitigate the effect of reflections at the connection point of the stub to the line. If the line impedance is about 50 ohms and the impedance of the stub is about 75 ohms, a resistor of approximately 20 - 25 ohms is typically chosen usually by trial and error for the best results under certain conditions. This resistor has the possibly undesirable effect of attenuating the voltage swing of the signal as the signal passes through the resistor, requiring a driver on the stub to have a proportionately larger voltage swing. Another undesirable effect is the RC delay due to the added series resistor and the device capacitance. The resistors and stubs also lead to low-fidelity signals at the devices. Also, as shown in Fig. 3, the line is terminated by resistors 350 at both ends to minimize reflections from the ends of the line. This requires that the drivers on the line supply twice as much steady state current as compared to a line terminated at only one end.

As mentioned above, the need to incorporate memory modules into the design of the modular system may also imply the use of connectors. In general, connectors have undesirable characteristics for operating at high frequency, such as inductance, capacitance, or crosstalk which introduces noise from one line into another line. Failure to take the connector characteristics into account leads to non-uniform arrival characteristics and low-fidelity signals when crossing a connector boundary resulting in lower performance (due to longer settling times, reduced noise margin or different signal propagation speed) from the modular system using the lines.

The physical shape, size and construction of the memory module is important to consider as well. The physical nature of the memory module may force the IC devices mounted on the module to be arranged in a less-than-optimal topology for the high frequency transmission line layout. High frequency signaling typically requires that electrical paths be controlled; signal delays need to be minimized or matched

and impedance needs to be tightly controlled for high frequency operation, where high frequency means frequencies in the range of 200 megaHertz to at least 1,000 megaHertz.

For the foregoing reasons, there is a need for a bus connecting to a plurality of devices which has uniform arrival times of high-fidelity signals to the devices on the bus, even when modules and connectors are employed to build a computer subsystem in which the bus is used and despite the physical size, shape and construction of the module and the number of devices mounted on it.

SUMMARY OF THE INVENTION

The present invention is directed to a high frequency bus system which insures uniform arrival times of high-fidelity signals to the devices on the high frequency bus, despite the use of the bus on modules and connectors. A high frequency bus system includes a first bus segment having one or more devices connected between a first and a second end. The first bus segment has at least a pair of transmission lines for propagating high frequency signals and the devices are coupled to the pair of transmission lines. The high frequency bus system also includes a second bus segment which has no devices connected to it. The second bus segment also has at least a pair of transmission lines for propagating high frequency signals. The first end of the first segment and second end of the second segment are coupled in series to form a chain of segments and when two signals are introduced to the first end of the second bus segment at the substantially the same time, they arrive at each device connected to the first bus segment at substantially the same time. Also, when two signals originate at a device connected to the first bus segment at substantially the same time, they arrive at the first end of the second bus segment at substantially the same time. Uniform arrival times hold despite the use of connectors to couple the segments together, despite the segments being located on modules, without the need for stubs, despite the presence of routing turns in the

segments and despite the type of information, such as address, data or control, carried by the signals.

In a preferred embodiment, the bus runs along a motherboard (second segment), onto one end of a memory module, then along the memory module (first segment), and exits the other end of the memory module along another motherboard segment to a next memory module. By running the bus through the module, stubs between a motherboard bus and each device are avoided, eliminating the need for resistors to compensate for reflections caused by the stubs. Preferably, each bus goes to all devices, using a control packet to select particular devices as needed.

In a preferred embodiment, the uniform arrival time is also insured by a number of routing and impedance modifying techniques. In particular, the total length travelled by different busses is equalized by having shorter horizontal length bus lines connected with corresponding longer vertical lengths joined by a right angle. In another aspect, parallel bus lines are used to equalize impedance for internally routed lines, since they have dielectric material on both sides, compared to a surface line with dielectric on only one side.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 shows the routing topology for a prior art modular memory system.

Fig. 2 shows a prior art memory module of the type employed in Fig. 1.

Fig. 3 shows a circuit model of a transmission line of a prior art module used in Fig. 1, including the terminators and the stubs.

Fig. 4 gives a perspective view of an embodiment of the present invention.

Fig. 5 shows a representative module of an embodiment of the present invention.

Fig. 6 gives a pictorial view of the routing of the transmission lines for an embodiment of the present invention

Fig. 7 gives a pictorial view of the routing of the transmission lines on the motherboard for an embodiment of the present invention.

Fig. 8 gives a simplified view of the routing of a data transmission line for an embodiment of the present invention.

Fig. 9 shows a simplified view of the routing of a clock transmission line for an embodiment of the present invention

Fig. 10 shows a simplified view of an alternative routing of a data transmission line for an embodiment of the present invention.

Fig. 11 shows another alternative routing for a data transmission line of an embodiment of the present invention.

Fig. 12 depicts a circuit model of a data transmission line near the connection point between the motherboard and the module for an embodiment of the present invention.

Fig. 13A depicts a circuit model of a data transmission line on the module for an embodiment of the present invention.

Fig. 13B depicts a circuit model of a device load on the module for an embodiment of the present invention.

Fig. 14A shows the routing of the transmission lines on the module near the right angle turn of the lines for an embodiment of the present invention.

Fig. 14B shows a magnified view of the routing of the transmission lines on the module near the right angle turn for an embodiment of the present invention.

Fig. 15A shows a perspective view from one side of the module of the routing of the transmission lines on the module near the right angle turn of the lines for an embodiment of the present invention.

Fig. 15B shows a perspective view from the opposite side shown in Fig. 15A of the module of the routing of the transmission lines on the module near the right angle turn of the lines for an embodiment of the present invention.

Fig. 16 shows a cross-section view of the module construction for an embodiment of the present invention.

09839768-041901

Fig. 17 shows a microstrip with the important parameters for computing the impedance and velocity of the signals on the transmission line of a module for an embodiment of the present invention.

5 Fig. 18 shows a stripline with the important parameters for computing the impedance and velocity of the signals on the transmission line for an embodiment of the present invention.

10 Fig. 19 shows a stripline with two internal transmission lines and the important parameters for computing the impedance and velocity of the signals on the transmission line for an embodiment of the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENT

15 Fig. 4 shows a perspective view of an embodiment of the present invention. Motherboard 410 acts as the substrate to which device modules 420 are coupled. A device 400 on the motherboard acts as a master which controls the devices on the modules. A high frequency bus 450 is routed from master 400
20 on the motherboard, through each module and finally to a terminator 440 on the motherboard. Each module 420 has a marking or key 430 to assure that each module is coupled with the same orientation to the motherboard. A device module 420 may or may not have devices mounted on it. A module with no
25 devices still has the bus routed through it and is termed a continuity module. Continuity modules are used to preserve the continuity of the bus when modules must be removed.

A more detailed view of module 420 is shown in Fig. 5 as module 520. Module 520 has devices 510 mounted one side
30 or both sides of the module. The module also includes a set of edge fingers 540, 550 for coupling to the motherboard and bus transmission lines 530 which enter the module on the edge fingers 540 of one side, make a right angle turn at a feedthrough hole 580, traverse along the length of the module
35 connecting to the devices on either side of the module, make another right angle turn and exit on the opposite side edge fingers 550. Edge fingers have a certain distance between their center lines (referred to as pitch) the connection

05839768 "041901
TOP40" 8926860

points on the devices have a certain pitch. The pitch puts an upper limit on the width of a line connecting to the fingers 540, 550 or the devices. For example, if the distance between edge fingers (measured from center line to center line) is about 30 mils (thousands of an inch), then the lines 530 which connect to the edge fingers must be less than about 22 mils in width, allowing for an 8 mil separation between lines. Also depicted are edge fingers 595 connecting to a reference plane within the module by means of lines 585. This preserves the relationship of grounds between signals as the signals pass through the connector joining the module to the motherboard.

Notches 560 on the side having the edge fingers are used as a key to assure the proper orientation and electrical and parametric compatibility of the module when coupled to the motherboard. Notches 570 may be used by a clamping or retaining device to hold the device module in place on the motherboard.

Fig. 6 gives a pictorial view of the routing of a data transmission line and a clock transmission line for an embodiment of the present invention. Data line or lines 650 are routed on the motherboard from master or controller 600 to device module 620a. The lines then traverse the length of the module substantially parallel to the long side of the module and exit at the opposite end from the entry onto the module. Along the length of the line on the module each device 655, situated on either side of the module, is connected to the line. The line is routed on the motherboard again until it enters the next module 620b at one end, again traversing the length of the module substantially parallel to the long side of the module and exiting at the opposite end. At the exit from module 620c the line then runs on the motherboard to a termination device 640. Clock line 660 starts at a clock source 680 and enters module 620c. It traverses the module 620c connecting to each device and exits from 620c onto the motherboard. The clock line 660 continues similarly until it reaches controller 600. Near or within controller 600 the clock line 660 is looped back to join to clock line 670 which travels back through the modules in a similar fashion to the

routing for the data line. Upon exiting the module 620c the clock line runs on the motherboard to a termination device 645. It is important for the proper operation of the devices that the clock lines 660 and 670 be routed in a fashion similar to the data line 650 and that no relative changes in direction of propagation occur between the two lines. This means that if a signal is propagating on the data line 650 of module 620a away from the dot (the orientation dot shown in the lower left of each module) then the signal on the clock line 670 propagates in the same direction and the signal on clock line 660 propagates in the opposite direction. This relationship is preserved between the data line and the clock lines on the motherboard and on the other modules.

Fig. 8 shows a more simple version of the routing of a data line in an embodiment of the present invention. Sections 800 are bus segments of the data line on the motherboard. Sections 810a-c are bus segments of the data line on the module with the module orientation key 830 also shown for the module. Couplings between the motherboard bus segments 800 and the module bus segments 850 are also shown. These couplings can be either permanent connections such as connections by means of pins on the module soldered into through-holes on the mother board, or removable connections such as in the case of a connector. The terminator for the bus segment 840 is shown at the end of the line. Also depicted are right angle turns 860 at both ends of the bus segment on the module. These right angle turns are necessary to assure substantially equal lengths for each data line routed according to Fig. 8, when lines connect to the module from a common edge, and is discussed below. Note also that there is symmetry to the routing of the lines on the module. In particular, the right angle turns are on both sides of the device locations. This allows signals to enter and leave the module from either end relative to the orientation key. In either case, the signals still arrive at the same time at the device locations if they started at the same time at the non-terminated end.

In Fig. 8 it is allowed that bus segments 810a-c can be segments on the mother board if there are no modules. These segments 810a-c are distinguished from segments 800 by whether there are devices present on the segment. Sections 800 have no devices connected between the ends of the segment whereas sections 810a-c have devices connected between the ends.

The routing of the clock segments is shown in Fig. 9. Again segments 900 are routed on the motherboard and segments 900a-c are those routed on the module. In the case of the clock line, the clock originates from clock source 980 typically located on the motherboard is routed through the modules to the loop 905, typically on the motherboard, back through the modules 900a-c and terminates at the termination device 940, also typically on the motherboard. Also shown are the right angle turns 960 that the clock lines must make on the modules 910a-c. Keys 930 are shown to indicate that all modules are coupled to the motherboard with the same orientation. Again, the so-called module segments 910a-c need not be physically separate from the motherboard, but instead may be segments on the motherboard having devices connected between the ends of the segments. However, no devices are connected between the ends of segments 900.

Fig. 10 shows an alternate embodiment of the routing of the data line. In the case of Fig. 10 the routing of the line terminates at a termination device 1040 on the module 1010 rather than the motherboard. Again the keys 1030 are shown to give the orientation. Thus, the module having the termination device is the last module in the chain of segments. The clock segment of Fig. 9 could also have the termination device on the last module in a manner similar to the segment of Fig. 10.

Fig. 11 shows the routing of a data line according to an embodiment of the present invention such that the route always enters the keyed end of the module 1110 requiring two folds 1100 on the motherboard. The difference between this embodiment and the one shown in Figs. 8 and 9 is that the routes in the latter embodiment are shorter, thus reducing the

time it takes a signal to propagate from one end to the other end of the line. In some cases it may be necessary to route the bus segments in the chain as shown in Fig. 11.

When connectors are used as the means of coupling between bus segments, certain characteristics of the connector are important with respect to their effect on the transmission lines coupled by the connector. Fig. 12 shows a circuit model of a connector and the transmission lines on either side and near the connection made by the connector. Section 1240 shows an equivalent circuit representation of the transmission line segment on the motherboard near the connector. Section 1250 shows the segment of transmission line on the module near the connector and section 1260 shows a circuit model of the connector itself. In the motherboard section 1240, the unloaded impedance of the line 1200 is Z_0 , which is preferably in the range of 22 to 32 ohms and more preferably 28 ohms. Motherboard section 1240 also shows the capacitance CMB by capacitor 1220. This capacitance results from a through-hole, surface mount pad or other capacitive structure, placed on the mother board in the path of the line 1270. The connector section 1260 includes the capacitance CPIN of the connector pin shown as capacitor 1240 and the inductance LPIN of the pin shown as inductor 1230. Finally, section 1250 includes a pad on the module having capacitance CPAD and shown as capacitor 1255. Line 1210 in module section 1250 has the same preferred impedance Z_0 as the line 1200 in motherboard module 1240. By controlling the value of CMB and CPAD, the equivalent impedance of the pin is altered to become the preferred impedance Z_0 . It is preferred for connectors used in connection with an embodiment of the present invention that the pin inductance be in the range of 2 to 3 nanohenrys (Nh). A typical value is approximately 2.3 Nh. The pin capacitance is in the range of 0.5 to 1.0 pico-farads (Pf) and is typically about 0.6 Pf. If the value of CMB and CPAD are made to be about 1 Pf, the total capacitance near the connector is about 2.6 Pf and the total inductance is about 2.3 nH. The result is an effective impedance of about 28-30 ohms for the connector. Thus, a signal propagating from the motherboard

section 1240 through the connector section 1260 on onto the module section 1250 encounters no significant change of impedance and no significant reflection is generated. Furthermore, the time to pass through the connector, i.e., from beginning of section 1240 to the end of section 1250 is uniform across all of the signals.

Fig. 13A shows a circuit model for the transmission line segment on a module in an embodiment of the present invention, where the term "module" implies the presence of devices connected to the transmission line segment. In this figure, connection points 1350 represent the coupling between the motherboard and the module, which may be by means of a connector whose characteristics were discussed above. Line sections 1360 are those portions of the transmission line having impedance Z_0 and no device connections, and line sections 1320 comprise sections of transmission line 1355 having an unloaded line impedance of Z_1 and a device load 1340. Fig. 13B shows an equivalent circuit for the device load indicating that it can be modeled approximately as a capacitor CDL. The value of CDL is approximately in the range of 2 - 3 Pf, but preferably about 2.5 Pf for a single device connection. To avoid discontinuities between the Z_0 section of line 1360 and the portion of line 1320 having the device load, the Z_1 impedance of line sections 1355 is increased substantially over the Z_0 value so that the equivalent impedance of line section 1320 matches that of the Z_0 section 1360. If the nominal impedance of the Z_0 section 1360 is about 28 ohms and a device load is to a first order capacitive and approximately 2.5 pF, then Z_1 is made approximately 70 ohms. The result is that the additional capacitance reduces the impedance of section 1320 to about 28 ohms according to the well known relationship that the loaded impedance Z_1' equals the unloaded impedance Z_1 multiplied by the square root of the ratio of the unloaded capacitance C_1 to the total capacitance CT .

$$Z_1' = Z_1 \sqrt{C_1/CT} \text{ where } CT = C_1 + CDL,$$
where CDL is the device load.

Thus, adding capacitance to the line lowers its impedance. A constraint on the above calculation is that the length of section 1320 must be on the order of an inch or less, preferably about 0.25 inches, implying that the device loads must be spaced by about 0.25 inches apart. If the device loads are spaced more widely, say by greater than 1 inch, then a line with an impedance of Z_0 must be placed between the device loads so that distance over which the device loading has an effect is reduced to less than an inch.

The reason for the distance limitation is based the transition time of the signal and how far that signal can travel during its transition time. Preferably, the signal to be propagated over the line has a rise time of about 200 pico-seconds (pS). During that time the signal will travel about an inch for a line on the surface of the module, assuming that signals on the surface of the module have a flight time of 150 to 200 pS/inch (to be discussed below). If there are no device loads on the module because no devices are present (in the case of a continuity module) then line sections 1320 are not present. Only line sections 1360 having an impedance of Z_0 are present and run the length of the module.

Fig. 14A shows more detail regarding the routing of the lines on the module 1400 and, in particular, shows the right angle turn taken by a line 1420 routed on the near surface of the module and a line 1430 routed on the far surface of the module. The purpose of the right angle turn is to assure equal and shortest lengths of lines across the module. This helps to assure uniform arrival times of the signals at each of the devices on the module. These lines 1420 and 1430 are received by vias 1460 at which point the lines change direction by 90 degrees to run parallel with the length of the module on a pair of internal sections 1440a,b per line. Dual internal tracks are necessary to match the impedance of the line on the surface 1320 to the impedance of the line internally 1440a,b as will be discussed below. Thus, sections 1420 and 1440a,b comprise section 1360 in Fig. 13A. Sections 1440a,b continue their routes down the length of the module until reaching the section where the devices are

located on the module. At this point the sections come back to the surface of the module by means of vias 1410 to connect to the device. Sections 1470 of the line comprise section 1320 in Fig. 13A. Sections 1470 continue along the length of the module until a point is reached where there are no more devices to connect to the line. The lines then exit the board in a fashion similar to how the lines entered, again by dual tracks internally for additional routing along the length, then into vias to make the right angle turn and then running on either surface of the module to the connection point of the module with the motherboard. Continuity modules placed anywhere except after the last module having devices on it, should have equal length lines to assure that signals entering the module at the same time leave the module at the same time. Therefore, if turns are used to route the lines of the continuity module, they should be right angle turns to maintain equal lengths for the lines.

Fig. 14B shows a magnified view of the right angle routing. As can be seen from the figure the length of any line from point A to point B is always equal to 10 units. Line 1420 is routed on one side of the module to a feed through hole 1460 at which point the line 1440a continues on a different layer. Line 1430 is routed on the opposite side of the module from line 1420 to a feed through hole 1460 at which point the line 1440b continues on the same layer as 1440a. This right angle routing thus requires at least three signal layers for routing if lines 1420 and 1430 must be routed on opposite sides of the module, possibly to match the pitch of any connector to which they may connect. Only two layers are required when lines 1420 and 1430 are routed from the same side of the module.

Figs. 15A and 15B give a perspective view of the sections of line and the internal construction of the module. In Fig. 15A sections 1520 correspond to sections 1420 in Fig. 14 and sections 1540a,b correspond to sections 1440a,b in Fig. 14. Sections 1510 in Fig. 15 correspond to sections 1470 in Fig. 14A and are the sections that run at the site of the devices on the module. Fig. 15A also shows the two reference

planes 1560 internal to the module to maintain a symmetric structure. These planes are typically at ground potential and provide a return path for the currents travelling on the lines. Vias 1550 connect the two reference planes 1560 together at various points near the via 1525 and internal section 1575 which connects the lines 1520 to the dual internal sections 1540a,b.

Fig. 15B shows the two reference planes 1560, a section 1522 entering on the opposite side of the module from section 1520 in Fig. 15A. Internal sections 1540a,b are shown as in Fig. 15A between the two reference planes 1560. Section 1570 of the line in Fig. 15B results from the plating through of via 1525 in Fig. 15A. The plating is done so that the dual internal tracks 1540a,b can be connected to section 1520 in Fig. 15A. Filled via 1555 is the result of plating through via 1550 in Fig. 15A which serves to connect the two reference planes together. Since section 1520 in Fig. 15A only connects on one side of the module but couples to a pair of internal sections 1540a,b, it is necessary to provide for a return path for the image currents induced by the currents on the internal sections. The section 1565 in Fig. 15A connecting the two reference planes together provides the return path for the image currents. The section 1565 connecting the two reference planes also has the purpose of reducing cross-talk between the signal lines. This effect results from the placement of the section 1565 near the section 1575 coupling the two internal sections to each other and to the signal section 1520, causing the section 1565 to act as a partial shield between the signal lines.

Fig. 16 shows a cross-section view of the internal construction of the module for an embodiment of the present invention. As in Fig. 15A and 15B, dual internal sections 1640a and 1640b, shown in View A, are shown between the two reference planes 1660. Section 1620 is a line section running on the surface of the module away from the sites of the devices and corresponds to section 1520 in Fig. 15A. Section 1675, in Views A and B, is the section that connects the dual internal sections 1640a and 1640b together. Section 1625 is

the section of the line that runs on the surface of the module near the location of the IC devices, leaving the space between the two reference planes open for additional lines or reference planes. The additional planes or lines 1670a,b, are shown in View D of Fig. 16. In View D the additional planes may be a plane which provides power to the devices (a power plane) and a reference voltage (Vref) plane for use by all of the devices. The additional planes are placed close to the reference plane to improve the quality of coupling to the reference planes. This helps reduce noise on the additional planes. In another embodiment, shown in View C, section 1627, runs internally rather than on the surface near the location of the IC devices. View E of Fig. 16 shows an alternative embodiment for routing lines internally. Rather than route the lines as in View A where the two lines 1640a and 1640b are connected to the same signal 1620 by feedthrough 1675 and run in parallel between the same reference planes 1660, View E shows a case where the two lines 1640a and 1640b are run in parallel but have their own reference planes. In particular, line 1642a has 1661 and 1662 as its reference planes, and 1642b has 1664 and 1665 as its reference planes. In View E, segments of line having IC devices connected between the ends have one of the lines, either 1640a or 1640b, eliminated. In segments of line having no IC devices connected between the ends, the lines run in parallel connected to the same signal.

Because sections of the line run on the surface of the module and some sections run internally between the two reference planes, the sections of line have different characteristics such as impedance and time of travel (flight time). Sections of line, such as either 1640a or 1640b individually, have a different impedance than section 1620 or 1625 running on the surface of the module. These two types of lines are shown in Figs. 17 and 18. It is important to match the impedance of the two types of lines and to compensate for any difference in flight time to achieve uniform arrival times of signals at the devices.

Fig. 17 shows, in cross-section, the important parameters for the section of line running on the surface of

the module away from the presence of devices, i.e., section 1620 in Fig. 16. In Fig. 17, the line 1700 has width, W , thickness, t , and runs over a dielectric material 1710 having thickness, S , and permittivity, ϵ_r . Ground plane 1720 is located opposite the line 1700 on the other side of the dielectric. In this construction, commonly referred to as a microstrip, the impedance of the line is determined to a first order by the ratio W/S and the permittivity ϵ_r of the dielectric. (See Matick, Transmission Lines For Digital and Communication Networks, IEEE Press, 1995, p. 326.). It is preferred to have the width of the line, W , in the range of 20 - 24 mils, preferably 22 mils to match the pitch of any connector coupled to the line. To construct a line having an impedance of approximately 25-30 ohms and preferably 28 ohms with a material having a permittivity of 4.5, the separation of the line S from the ground plane is in the range of 4.4 to 5.3 mils, preferable 5 mils, giving a W/S ratio of about 4.5. A smaller width line W or a greater separation distance S results in an impedance that is greater than the preferred range of 25 to 30 ohms. Thus, segments of line travelling on the surface of the module and without devices connected between the ends have an impedance in the range of 25 to 30 ohms.

To construct a line of the type shown in Fig. 17, but having an impedance of approximately 70 ohms with the same dielectric material for lines on the surface of the module and to which the devices connect, the W/S ratio is altered to become approximately equal 1.0. Thus, by altering the width of the trace from approximately 22 mils to 5 mils, and holding the distance S constant, a line running on the surface of the module near the IC devices, section 1625 in Fig. 16, is constructed having an impedance of about 70 ohms. As discussed above, when this line is taken in combination with the device load, the impedance of the section is reduced to about 28 ohms, thus matching the impedance of an line section which is does not have any devices connected to it.

Fig. 18 shows, in cross-section, the important parameters for a section of line having no devices connected

09039760-04904
TOP SECRET

to it and running internally between the two reference planes of the module, i.e., section 1640a or 1640b in Fig. 16. In Fig. 18, the width of the line 1800 is W and its thickness is t . The separation distance between the two reference planes 1820 is S' and the dielectric filling the separation distance has a permittivity of ϵ_r . In this construction, the impedance of the line is determined to a first order by the W/S' ratio, the t/S' ratio and the permittivity ϵ_r . (See Matick, Transmission Lines For Digital and Communication Networks, IEEE Press, 1995, p. 327.) The W/S' ratio will typically be less than 1 for lines having practical values of trace widths, W , of about 22 mils or less, determined by connector pitch and a practical separation distance S' between planes of approximately 35 mils determined by module thickness required for insertion of the module into a socket. Under these circumstances, the impedance of a line shown in Fig. 18, commonly referred to as a stripline, will be approximately 50 ohms when W is approximately 14 mils, S' is approximately 35 mils and t is approximately 1.4 mils (for 1 ounce copper plating). This impedance does not match the preferred impedance of the other lines to which the internal lines connect assuming the use of the same dielectric material. Therefore, two lines, connected to the same signal, are run in parallel internally as is shown in Figs. 16 and 19. The use of two lines reduces the impedance from 50 ohms to about 28 ohms (rather than 25 ohms) due to inductive coupling that is present between the lines. Thus, the use of the dual parallel lines creates a internal section of line that matches the impedance of the other lines to which the internal section connects.

To construct a line of the type in Fig. 18 having an impedance of approximately 70 ohms for sections having devices connected, such as section 1627 in View C of Fig. 16, the W/S' ratio is altered to about 0.15 by decreasing the width of the trace from about 14 mils to 5.5 mils. The thickness of the trace, t , and the other parameters are held constant. To construct a line of the type in Fig. 18 having an impedance of approximately 50 ohms, such as for section 1642a or 1642b in

View E of Fig. 16, the separation distance S' is about 15 mils, the line width W is about 3 mils, and the line thickness is 1.4 mils. This creates a W/S' ratio of about 0.2 and a t/S' ratio of about 0.1. The impedance then is about 56 ohms for ϵ_r equal to 4.5. Thus, two lines 1642a and 1642b are run in parallel and connected to the same signal to give an impedance of 28 ohms for segments of the line without IC devices. Near the IC devices one of the lines 1642a or 1642b is eliminated resulting in a segment of line having an impedance of 56 ohms for segments of the line with IC devices. With device loads taken into account, the impedance is reduced to 28 ohms.

There is still a problem associated with the use of the two types of lines as shown in Figs. 17 and 18, even if the lines are constructed so that the impedance is the same for each type. The problem is that the time for a signal to travel a unit distance (flight time) for a line constructed as in Fig. 17 is faster than the flight time for a line constructed as in Fig. 18. The expression for the flight time TF is

$$TF = TC \sqrt{\epsilon_r},$$

where TC is the flight time of a signal in free space (approximately 84.7 picoseconds per inch = 3.33 pS/millimeter) and ϵ_r is the effective relative permittivity of the medium in which the signal propagates. For lines constructed according to Figs. 18 or 19, the flight time is about 180 pS/in (picoseconds per inch), because the effective relative permittivity is about 4.5. The flight time of a line constructed according to Fig. 17 is more difficult to determine because the effective relative permittivity has a contribution from both the dielectric 1710 and air surrounding the top and sides of the conductor 1700. The effective relative permittivity is less than that of the dielectric and in the range of 3.0 to 3.15. Thus, according to the equation, flight time for such a line is in the range of about 145 pS/in to 150 pS/in. Compared to the line in Fig. 18, the line in Fig. 17 is about 17% to 20% faster.

To assure that signals travelling on the lines of the module arrive at the IC device at the same time requires some form of compensation which takes into account the different flight times of the different types of lines. In one embodiment of the present invention, the physical line is first made to have the same length by routing the lines with a right angle turn as shown on Fig. 14A,B. Right angle turning or "folding" of the line assures that the length of section 1420 added to the length of section 1440a is the same for all of the folded lines. This helps to cut down on the differences between arrival times at the IC device locations 1410 in Fig. 14A,B. Next, the slowest path to a device is identified. In the embodiment shown in Fig. 14A,B, the slowest path is the path comprising sections 1420 and 1440a. This path is slowest because the signals on this path travel the longest distance over a section of the line with the slowest flight time. Finally, delay matching segments, 1480 in Fig. 14 (and 530 in Fig. 5) are added to the faster paths so that the signals on them arrive at the IC device site 1410 at the same time as the slowest path. Thus, folding the group of lines, assures that the lines have equal lengths except for the delay compensating sections which are then added to compensate for the unequal flight times between the two types of lines. After compensation, signals entering the module take the same time to travel to the site of the IC devices 1410 regardless of the particular routing path of the line.

Continuity modules used to couple the bus to modules having devices on them, should have delay matching segments added if the routes on the continuity module use right angle turns with portions of the turn having different flight times. The delay matching segments assure that signals entering the continuity module at the same time, leave the continuity module at the same time. A continuity module connected to the termination device need not have any delay matching segments added as uniform arrival times of signals at the termination device is not necessary.

Although the invention has been described in considerable detail with reference to certain embodiments

thereof, other embodiments are possible. For example, Fig. 7 shows another embodiment of the present invention in which the devices are mounted on the motherboard directly. The bus has sections 790 in which no devices are present and a section 795 with devices connected. The data lines are routed from the controller or master 700 to the individual devices 755 and then to the terminator device 740. Near the devices, the width of the data lines is altered to maintain an impedance that matches the nominal impedance of the sections 790 where there are no devices 755. Where data lines must turn in their route from the controller to the devices or from the devices to the terminating device, the turns are right angle turns to maintain the same length of each line. In addition, delay matching segments are added where there are sections of the route that have different flight times. Thus the arrival time of signals to the devices in the device segment of the line 795 is still assured.

In Fig. 7, the clock lines are routed from the clock source 780 to the devices 755 to the controller 700. At or near the controller the clock line reverses direction by means of clock loop 725 and continues on line 770 to a terminator device 745. The clock lines have uniform impedance throughout the length of the line, again by altering the width of the clock sections near the devices to match the impedance of sections 790 without devices. Again, right angle turns are employed to maintain the same length of line as the data lines. Delay matching segments are added where there are sections of the route that have different flight times. Another embodiment that is possible is locating the controller 400 depicted in Fig. 4 on a module 420 rather than on the motherboard 410. Therefore, the spirit and scope of the appended claims should not be limited to the description of the preferred versions contained therein.